

Search: The ACM Digital Library The Guide

+metadata +location +lock +portion



THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before March 1999 Terms used metadata location lock portion

Found 68 of 94,677

Sort results by

Display

results

relevance

expanded form

Save results to a Binder 3 Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 1 - 20 of 68

· Result page: **1** <u>2</u> <u>3</u> <u>4</u>

Relevance scale 🔲 📟 📟 📟

Frangipani: a scalable distributed file system

Chandramohan A. Thekkath, Timothy Mann, Edward K. Lee

window

October 1997 ACM SIGOPS Operating Systems Review, Proceedings of the sixteenth ACM symposium on Operating systems principles, Volume 31 Issue 5

Full text available: pdf(2.20 MB)

Additional Information: full citation, references, citings, index terms

2 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997 Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research

Full text available: pdf(4.21 MB)

Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

3 HFS: a performance-oriented flexible file system based on building-block compositions. Orran Krieger, Michael Stumm



August 1997 ACM Transactions on Computer Systems (TOCS), Volume 15 Issue 3

Full text available: pdf(383.87 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

The Hurricane File System (HFS) is designed for (potentially large-scale) shared-memory multiprocessors. Its architecture is based on the principle that, in order to maximize performance for applications with diverse requirements, a file system must support a wide variety of file structures, file system policies, and I/O interfaces. Files in HFS are implemented using simple building blocks composed in potentially complex ways. This approach yields great flexibility, allowing an application ...

Keywords: customization, data partitioning, data replication, flexibility, parallel computing, parallel file system

Serverless network file systems

Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, Randolph Y. Wang

February 1996 ACM Transactions on Computer Systems (TOCS), Volume 14 Issue 1

Full text available: mpdf(2.69 MB)

Additional Information: full citation, abstract, references, citings, index

We propose a new paradigm for network file system design: serverless network file systems. While traditional network file systems rely on a central server machine, a serverless system utilizes workstations cooperating as peers to provide all file system services. Any machine in the system can store, cache, or control any block of data. Our approach uses this location independence, in combination with fast local area networks, to provide better performance and scalability th ...

Keywords: RAID, log cleaning, log structured, log-based striping, logging, redundant data storage, scalable performance

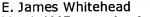
Serverless network file systems

T. E. Anderson, M. D. Dahlin, J. M. Neefe, D. A. Patterson, D. S. Roselli, R. Y. Wang December 1995 ACM SIGOPS Operating Systems Review, Proceedings of the fifteenth ACM symposium on Operating systems principles, Volume 29 Issue 5

Full text available: pdf(2.48 MB)

Additional Information: full citation, references, citings, index terms

World Wide Web distributed authoring and versioning (WebDAV): an introduction



March 1997 StandardView, Volume 5 Issue 1

Full text available: pdf(81.76 KB)

Additional Information: full citation, references, citings, index terms, review

7 The Integrated Dictionary/Directory System

Frank W. Allen, Mary E. S. Loomis, Michael V. Mannino

June 1982 ACM Computing Surveys (CSUR), Volume 14 Issue 2

Full text available: pol(2.71 MB)

Additional Information: full citation, references, citings, index terms

8 Workshop on compositional software architectures: workshop report

May 1998 ACM SIGSOFT Software Engineering Notes, Volume 23 Issue 3

Full text available: pdf(2.91 MB) Additional Information: full citation, index terms

Application performance and flexibility on exokernel systems

M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Héctor M. Briceño, Russell Hunt, David Mazières, Thomas Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie October 1997 ACM SIGOPS Operating Systems Review, Proceedings of the sixteenth ACM symposium on Operating systems principles, Volume 31 Issue 5

Full text available: Report (2.39 MB)

Additional Information: full citation, references, citings, index terms



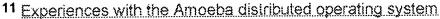
¹⁰ Making sense of software engineering environment framework standards

Barbara Cuthill

December 1994 StandardView, Volume 2 Issue 4

Full text available: pdf(1.67 MB)

Additional Information: full citation, references, index terms



Andrew S. Tanenbaum, Robbert van Renesse, Hans van Staveren, Gregory J. Sharp, Sape J.

December 1990 Communications of the ACM, Volume 33 Issue 12

Full text available: pdf(2.71 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

The Amoeba project is a research effort aimed at understanding how to connect multiple computers in a seamless way [16, 17, 26, 27, 31]. The basic idea is to provide the users with the illusion of a single powerful timesharing system, when, in fact, the system is implemented on a collection of machines, potentially distributed among several countries. This research has led to the design and implementation of the Amoeba distributed operating system, which is being used as a prototype and veh ...

12 Requirements for distributed authoring and versioning on the World Wide Web

J. A. Slein, F. Vitali, E. J. Whitehead, D. G. Durand March 1997 StandardView, Volume 5 Issue 1

Full text available: pdf(96.62 KB)

Additional Information: full citation, references, citings, index terms

13 Recovery in the Calypso file system

Murthy Devarakonda, Bill Kish, Ajay Mohindra

August 1996 ACM Transactions on Computer Systems (TOCS), Volume 14 Issue 3

Full text available: pdf(318.88 KB)

Additional Information: fall citation, abstract, references, citings, index terms, review

This article presents the deign and implementation of the recovery scheme in Calypso. Calypso is a cluster-optimized, distributed file system for UNIX clusters. As in Sprite and AFS, Calypso servers are stateful and scale well to a large number of clients. The recovery scheme in Calypso is nondisruptive, meaning that open files remain open, client modified data are saved, and in-flight operations are properly handled across server recover. The scheme uses distributed state amount the client ...

Keywords: Calypso, cluster systems, distributed state, state reconstruction

14 The Zebra striped network file system

John H. Hartman, John K. Ousterhout

August 1995 ACM Transactions on Computer Systems (TOCS), Volume 13 Issue 3

Full text available: pdf(2.76 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Zebra is a network file system that increases throughput by striping the file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a log-structured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this ...







Keywords: RAID, log-based striping, log-structured file system, parity computation

15 Performance of cache coherence in stackable filing

J. Heidemann, G. Popek

December 1995 ACM SIGOPS Operating Systems Review, Proceedings of the fifteenth ACM symposium on Operating systems principles, Volume 29 Issue 5

Full text available: mpdf(2.00 MB)

Additional Information: full citation, references, index terms

16 Building a scaleable geo-spatial DBMS: technology, implementation, and evaluation Jignesh Patel, JieBing Yu, Navin Kabra, Kristin Tufte, Biswadeep Nag, Josef Burger, Nancy Hall, Karthikeyan Ramasamy, Roger Lueder, Curt Ellmann, Jim Kupsch, Shelly Guo, Johan Larson, David De Witt, Jeffrey Naughton

June 1997 ACM SIGMOD Record, Proceedings of the 1997 ACM SIGMOD international conference on Management of data, Volume 26 Issue 2

Full text available: mpdf(1.58 MB)

Additional Information: full citation, abstract, references, citings, index terms

This paper presents a number of new techniques for parallelizing geo-spatial database systems and discusses their implementation in the Paradise object-relational database system. The effectiveness of these techniques is demonstrated using a variety of complex geo-spatial queries over a 120 GB global geo-spatial data set.

17 Shoring up persistent applications

Michael J. Carey, David J. DeWitt, Michael J. Franklin, Nancy E. Hall, Mark L. McAuliffe, Jeffrey F. Naughton, Daniel T. Schuh, Marvin H. Solomon, C. K. Tan, Odysseas G. Tsatalos, Seth J. White, Michael J. Zwilling

May 1994 ACM SIGMOD Record, Proceedings of the 1994 ACM SIGMOD international conference on Management of data, Volume 23 Issue 2

Full text available: pdf(1,40 MB)

Additional Information: full citation, abstract, references, citings, index terms

SHORE (Scalable Heterogeneous Object REpository) is a persistent object system under development at the University of Wisconsin. SHORE represents a merger of object-oriented database and file system technologies. In this paper we give the goals and motivation for SHORE, and describe how SHORE provides features of both technologies. We also describe some novel aspects of the SHORE architecture, including a symmetric peer-to-peer server architecture, server customization through an extensible ...

18 An introduction to data warehousing: what are the implications for the network? Katherine Jones

February 1998 International Journal of Network Management, Volume 8 Issue 1

Full text available: pdf(145.35 KB) Additional Information: full citation, abstract, references, index terms

Data warehousing is an information systems environment, rather than a product. It has emerged as an essential business entity for sophisticated analysis of data. This article presents a clear overview of the implications of data warehousing for business. © 1998 John Wiley & Sons, Ltd.

19 The Zebra striped network file system

John H. Hartman, John K. Ousterhout

December 1993 ACM SIGOPS Operating Systems Review, Proceedings of the fourteenth ACM symposium on Operating systems principles, Volume 27 Issue 5

Additional Information:

Full text available: (1.93 MB)

full citation, abstract, references, citings, index terms

Zebra is a network file system that increases throughput by striping file data across multiple servers. Rather than striping each file separately, Zebra forms all the new data from each client into a single stream, which it then stripes using an approach similar to a logstructured file system. This provides high performance for writes of small files as well as for reads and writes of large files. Zebra also writes parity information in each stripe in the style of RAID disk arrays; this increase ...

20 Efficient support for irregular applications on distributed-memory machines Shubhendu S. Mukherjee, Shamik D. Sharma, Mark D. Hill, James R. Larus, Anne Rogers, Joel



August 1995 ACM SIGPLAN Notices, Proceedings of the fifth ACM SIGPLAN symposium on Principles and practice of parallel programming, Volume 30 Issue 8

Full text available: 📆 odf(1,36 MB)

Additional Information: full citation, abstract, references, citings, index terms

Irregular computation problems underlie many important scientific applications. Although these problems are computationally expensive, and so would seem appropriate for parallel machines, their irregular and unpredictable run-time behavior makes this type of parallel program difficult to write and adversely affects run-time performance. This paper explores three issues—partitioning, mutual exclusion, and data transfer—crucial to the efficient execution of irregular pro ...

Results 1 - 20 of 68

Result page: **1** $\frac{2}{3}$ $\frac{4}{4}$ next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Search: • The ACM Digital Library C The Guide

+metadata +location +lock +portion



THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before March 1999 Terms used metadata location lock portion

Found **68** of **94,677**

Sort results by

results

Display

relevance

expanded form

Save results to a Binder 3 Search Tips Open results in a new

window

Try an Advanced Search Try this search in The ACM Guide

Results 21 - 40 of 68

Result page: previous 1 2 3 4 next

Relevance scale ...

21 Experience with transactions in QuickSilver

Frank Schmuck, Jim Wylie

September 1991 ACM SIGOPS Operating Systems Review, Proceedings of the thirteenth ACM symposium on Operating systems principles, Volume 25

Full text available: pdf(1.66 MB)

Additional Information: full citation, abstract, references, citings, index terms

All programs in the QuickSilver distributed system behave atomically with respect to their updates to permanent data. Operating system support for transactions provides the framework required to support this, as well as a mechanism that unifies reclamation of resources after failures or normal process termination. This paper evaluates the use of transactions for these purposes in a general purpose operating system and presents some of the lessons learned from our experience with a complet ...

22 A toolkit for the incremental implementation of heterogeneous database management systems



Pamela Drew, Roger King, Dennis Heimbigner

October 1992 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 1 Issue 2

Full text available: Resistance (2.42 MB)

Additional Information: full citation, abstract, references, citings

The integration of heterogeneous database environments is a difficult and complex task. The A la carte Framework addresses this complexity by providing a reusable and extensible architecture in which a set of heterogeneous database management systems can be integrated. The goal is to support incremental integration of existing database facilities into heterogeneous, interoperative, distributed systems. The Framework addresses the three main issues in heterogeneous systems integration. First, it ...

Keywords: database toolkits, extensible databases, heterogeneous databases, heterogeneous transaction management, incremental integration, open architectures, reconfigurable architectures

23 VM-based shared memory on low-latency, remote-memory-access networks Leonidas Kontothanassis, Galen Hunt, Robert Stets, Nikolaos Hardavellas, Michał Cierniak, Srinivasan Parthasarathy, Wagner Meira, Sandhya Dwarkadas, Michael Scott May 1997 ACM SIGARCH Computer Architecture News, Proceedings of the 24th



annual international symposium on Computer architecture, Volume 25 Issue 2

Full text available: pdf(1.96 MB)

Additional Information: full citation, abstract, references, citings, index terms

Recent technological advances have produced network interfaces that provide users with very low-latency access to the memory of remote machines. We examine the impact of such networks on the implementation and performance of software DSM. Specifically, we compare two DSM systems---Cashmere and TreadMarks---on a 32-processor DEC Alpha cluster connected by a Memory Channel network. Both Cashmere and TreadMarks use virtual memory to maintain coherence on pages, and both use lazy, multi-writer releas ...

24 Electronic colloquia: idea and practice

Jochen Bern, Christoph Meinel, Harald Sack

September 1998 Proceedings of the 16th annual international conference on Computer documentation

Full text available: cdf(764.10 KB) Additional Information: full offation, references, citings, index terms

Keywords: electronic colloquia, electronic colloquium on computational complexity, electronic journals, electronic submission to conferences, information technology, symposium on theoretical aspects in computer science

25 Remotely-sensed geophysical databases: experience and implications for generalized DBMS



Guy M. Lohman, Joseph C. Stoltzfus, Anita N. Benson, Michael D. Martin, Alfonso F. Cardenas May 1983 ACM SIGMOD Record, Proceedings of the 1983 ACM SIGMOD international conference on Management of data, Volume 13 Issue 4

Additional Information: full citation, abstract, references, citings Full text available: pdf(1.85 MB)

This paper presents the characteristics of scientific remotely-sensed databases that are relevant to --- and pose unique challenges for --- general-purpose database management systems (DBMSs). We describe a prototype system that integrates geophysical data and its metadata from both satellite and in situ sources, using a relational general-purpose DBMS to manage the catalog and observational data, and a video optical disk to archive images. Based upon our experience with this application, ...

26 Sun MPII/O: efficient I/O for parallel applications



Len Wisniewski, Brad Smisloff, Nils Nieuwejaar

January 1999 Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)

Full text available: ndf(138.57 KB) Additional Information: full citation, references, citings, index terms

27 Hyperform: a hypermedia system development environment Uffe K. Wiil, John J. Leggett



January 1997 ACM Transactions on Information Systems (TOIS), Volume 15 Issue 1

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(362.84 KB) terms, review

Development of hypermedia systems is a complex matter. The current trend toward open, extensible, and distributed multiuser hypermedia systems adds additional complexity to the development process. As a means of reducing this complexity, there has been an increasing interest in hyperbase management systems that allow hypermedia system developers to abstract from the intricacies and complexity of the hyperbase layer and fully attend to

application and user interface issues. Design, developme ...

Keywords: advanced hypermedia system architecture, extensible hyperbase management system, object-oriented extension language

28 The Vesta parallel file system

Peter F. Corbett, Dror G. Feitelson

August 1996 ACM Transactions on Computer Systems (TOCS), Volume 14 Issue 3

Full text available: nxif(649.08 KB)

Additional Information: full citation, abstract, references, citings, index terms, review

The Vesta parallel file system is designed to provide parallel file access to application programs running on multicomputers with parallel I/O subsystems. Vesta uses a new abstraction of files: a file is not a sequence of bytes, but rather it can be partitioned into multiple disjoint sequences that are accessed in parallel. The partitioning—which can also be changed dynamically—reduces the need for synchronization and coordination during the access. Some control over the layout ...

Keywords: data partitioning, parallel computing, parallel file system

29 File-system development with stackable layers

John S. Heidemann, Gerald J. Popek

February 1994 ACM Transactions on Computer Systems (TOCS), Volume 12 Issue 1

Full text available: pdf(2.16 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Filing services have experienced a number of innovations in recent years, but many of these promising ideas have failed to enter into broad use. One reason is that current filing environments present several barriers to new development. For example, file systems today typically stand alone instead of building on the work of others, and support of new filing services often requires changes that invalidate existing work. Stackable file-system design addresses these issues in severa ...

Keywords: composability, file system design, operating system structure, reuse

30 QuickStore: a high performance mapped object store

Seth J. White, David J. DeWitt

May 1994 ACM SIGMOD Record, Proceedings of the 1994 ACM SIGMOD international conference on Management of data, Volume 23 Issue 2

Full text available: pdf(1.73 MB)

Additional Information: full citation, abstract, references, citings, index terms

This paper presents, QuickStore, a memory-mapped storage system for persistent C++ built on top of the EXODUS Storage Manager. QuickStore provides fast access to in-memory objects by allowing application programs to access objects via normal virtual memory pointers. The paper also presents the results of a detailed performance study using the OO7 benchmark. The study compares the performance of QuickStore with the latest implementation of the E programming language. These systems exemplify ...

31 A federated architecture for information management

Dennis Heimbigner, Dennis McLeod

July 1985 ACM Transactions on Information Systems (TOIS), Volume 3 Issue 3

Full text available: 📆 cdf(2.19 MB)

Additional Information: full citation, abstract, references, citings, index

terms

An approach to the coordinated sharing and interchange of computerized information is described emphasizing partial, controlled sharing among autonomous databases. Office information systems provide a particularly appropriate context for this type of information sharing and exchange. A federated database architecture is described in which a collection of independent database systems are united into a loosely coupled federation in order to share and exchange information. A federation consist ...

32 Recovery management in QuickSilver

Rober Haskin, Yoni Malachi, Gregory Chan

February 1988 ACM Transactions on Computer Systems (TOCS), Volume 6 Issue 1

Full text available: pxif(2.21 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

This paper describes QuickSilver, developed at the IBM Almaden Research Center, which uses atomic transactions as a unified failure recovery mechanism for a client-server structured distributed system. Transactions allow failure atomicity for related activities at a single server or at a number of independent servers. Rather than bundling transaction management into a dedicated language or recoverable object manager, Quicksilver exposes the basic commit protocol and log rec ...

33 Lightweight recoverable virtual memory

M. Satyanarayanan, Henry H. Mashburn, Puneet Kumar, David C. Steere, James J. Kistler February 1994 ACM Transactions on Computer Systems (TOCS), Volume 12 Issue 1

Full text available: pdf(1.73 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Recoverable virtual memoryrefers to regions of a virtual address space on which transactional guarantees are offered. This article describes RVM, an efficient, portable, and easily used implementation of recoverable virtual memory for Unix environments. A unique characteristic of RVM is that it allows independent control over the transactional properties of atomicity, permanence, and serializability. This leads to considerable flexibility in the use of RVM, potentially enla ...

Keywords: Camelot, Coda, RVM, Unix, logging, paging, persistence, scalability, throughput, truncation

34 The Rio file cache: surviving operating system crashes

Peter M. Chen, Wee Teck Ng, Subhachandra Chandra, Christopher Aycock, Gurushankar Rajamani, David Lowell

September 1996 Proceedings of the seventh international conference on Architectural support for programming languages and operating systems, Volume 31,

Full text available: pdf(1.12 MB)

Additional Information: full citation, abstract, references, citings, index <u>ierms</u>

One of the fundamental limits to high-performance, high-reliability file systems is memory's vulnerability to system crashes. Because memory is viewed as unsafe, systems periodically write data back to disk. The extra disk traffic lowers performance, and the delay period before data is safe lowers reliability. The goal of the Rio (RAM I/O) file cache is to make ordinary main memory safe for persistent storage by enabling memory to survive operating system crashes. Reliable memory enables a syste ...

35 CRL: high-performance all-software distributed shared memory

K. L. Johnson, M. F. Kaashoek, D. A. Wallach



December 1995 ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles, Volume 29 Issue 5

Full text available: pdf(2.02 MB)

Additional Information: full citation, references, citings, index terms

36 Applications: A visual interface for synchronous collaboration and negotiated transactions

Lutz Wegner, Manfred Paul, Jens Thamm, Sven Thelemann May 1996 Proceedings of the workshop on Advanced visual interfaces

Full text available: 📆 pot(2.43 MB)

Additional Information: full citation, abstract, references

This paper introduces a visual interface for computer-supported cooperative work (CSCW). The interface is an extension of the editor interface of ESCHER, a prototype database system based on the extended non-first-normal-form data model. In ESCHER, the nested table approach is the paradigm for presenting data, where presenting includes browsing, editing and querying the database. Interaction is achieved by fingers generalising the well-known cursor concept. When several users a ...

37 File server scaling with network-attached secure disks

Garth A. Gibson, David F. Nagle, Khalil Amiri, Fay W. Chang, Eugene M. Feinberg, Howard Gobioff, Chen Lee, Berend Ozceri, Erik Riedel, David Rochberg, Jim Zelenka June 1997 **ACM SIGMETRICS Performance Evaluation Review, Proceedings of the**

June 1997 ACM SIGMETRICS Performance Evaluation Review, Proceedings of the 1997 ACM SIGMETRICS international conference on Measurement and modeling of computer systems, Volume 25 Issue 1

Full text available: 📆 pdf(1.77 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> ierms

By providing direct data transfer between storage and client, network-attached storage devices have the potential to improve scalability for existing distributed file systems (by removing the server as a bottleneck) and bandwidth for new parallel and distributed file systems (through network striping and more efficient data paths). Together, these advantages influence a large enough fraction of the storage market to make commodity network-attached storage feasible. Realizing the technology's ful ...

38 Multiversion divergence control of time fuzziness

Calton Pu, Miu K. Tsang, Kun-Lung Wu, Philip S. Yu

November 1994 Proceedings of the third international conference on Information and knowledge management

Full text available: pdf(980.75 KB) Additional Information: full citation, abstract, references, index terms

Epsilon Serializability (ESR) has been proposed to manage and control inconsistency in extending the classic transaction processing. ESR increases system concurrency by tolerating a bounded amount of inconsistency. In this paper, we present multiversion divergence control (mvDC) algorithms that support ESR with not only value but also time fuzziness in multiversion databases. Unlike value fuzziness, accumulating time fuzziness is semantically different. A s ...

39 Performance of database workloads on shared-memory systems with out-of-order processors

Parthasarathy Ranganathan, Kourosh Gharachorloo, Sarita V. Adve, Luiz André Barroso
October 1998 Proceedings of the eighth international conference on Architectural
support for programming languages and operating systems, Volume 33, 32
Issue 11, 5

Full text available: gdf(1.62 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Database applications such as online transaction processing (OLTP) and decision support systems (DSS) constitute the largest and fastest-growing segment of the market for multiprocessor servers. However, most current system designs have been optimized to perform well on scientific and engineering workloads. Given the radically different behavior of database workloads (especially OLTP), it is important to re-evaluate key system design decisions in the context of this important class of applicatio ...

40 Session 3: Conceptual-to-internal mappings in commercial database systems D. S. Batory

Full text available: pdf(985.08 KB) Additional Information: full citation, abstract, references, citings



April 1984 Proceedings of the 3rd ACM SIGACT-SIGMOD symposium on Principles of database systems

Developing a commercial or specialized database system is an exceedingly costly and time consuming undertaking. A goal of this research is to demonstrate that a significant portion of a DBMS's software, specifically the physical database component, can be developed automatically from a small set of specifications. In this paper, we explain a new method of modeling physical databases and show that it provides a framework for realizing such a software development technology. Unlike any existing met ...

Results 21 - 40 of 68

Result page: previous 1 2 3 4 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player



Search: The ACM Digital Library The Guide

+metadata +location +lock +portion

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before March 1999 Terms used metadata location lock portion

Found 68 of 94,677

Sort results

bν Display

results

relevance expanded form

window

Save results to a Binder 3 Search Tips Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 41 - 60 of 68

Result page: previous 1 2 3 4

Relevance scale 🗆 📟 📟

41 Retrieving and visualizing video

Boon-Lock Yeo, Minerva M. Yeung

December 1997 Communications of the ACM, Volume 40 Issue 12

Full text available: pdf(2.01 MB)

Additional Information: full citation, references, citings, index terms

42 Integrating a dynamic lexicon with a dynamic full-text retrieval system

Peter G. Anick, Rex A. Flynn

July 1993 Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval

Full text available: pdf(1.24 MB)

Additional Information: full citation, abstract, references, citings, index terms

There has been a great deal of interest within the Information Retrieval community in evaluating the use of linguistic knowledge to improve the indexing and searching of textual databases. Such systems must often employ a lexicon to store information about the words and phrases comprising the application's domain. Unlike a static lexicon, a dynamic lexicon raises practical concerns about the coordination between the state of the lexicon and IR indexing sche ...

43 MMLite: a highly componentized system architecture

Johannes Helander, Alessandro Forin

September 1998 Proceedings of the 8th ACM SIGOPS European workshop on Support for composing distributed applications

Full text available: pdf(1.01 MB)

Additional Information: full citation, citings, index terms

44 Object orientation in multidatabase systems

Evaggelia Pitoura, Omran Bukhres, Ahmed Elmagarmid

June 1995 ACM Computing Surveys (CSUR), Volume 27 Issue 2

Full text available: 📆 pdf(4.85 MB)

Additional Information: full ciliation, abstract, references, citings, index terms, review

A multidatabase system (MDBS) is a confederation of preexisting distributed, heterogeneous, and autonomous database systems. There has been a recent proliferation of research suggesting the application of object-oriented techniques to facilitate the complex task of designing and implementing MDBSs. Although this approach seems promising, the lack of a general framework impedes any further development. The goal of this paper is to provide a concrete analysis and categorization of the various ...

Keywords: distributed objects, federated databases, integration, multidatabases, views

45 Interoperability for digital libraries worldwide

Andreas Paepcke, Chen-Chuan K. Chang, Terry Winograd, Héctor García-Molina April 1998 Communications of the ACM, Volume 41 Issue 4

Full text available: ndf(299.48 KB) Additional Information: full citation, references, citings, index terms

46 Free transactions with Rio Vista

David E. Lowell, Peter M. Chen

October 1997 ACM SIGOPS Operating Systems Review, Proceedings of the sixteenth ACM symposium on Operating systems principles, Volume 31 Issue 5

Full text available: 📆 odf(1.13 MB) Additional Information: full citation, references, citings, index terms

47 Versioning a full-text information retrieval system

Peter G. Anick, Rex A. Flynn

June 1992 Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval

Full text available: pdf(1.53 MB)

Additional Information: full cilation, abstract, references, citings, index

In this paper, we present an approach to the incorporation of object versioning into a distributed full-text information retrieval system. We propose an implementation based on "partially versioned" index sets, arguing that its space overhead and query-time performance make it suitable for full-text IR, with its heavy dependence on inverted indexing. We develop algorithms for computing both historical queries and time range queries and show how these algorithms can be applied to ...

48 Managing multimedia information in database systems

William I. Grosky

December 1997 Communications of the ACM, Volume 40 Issue 12

Full text available: pdf(1.91 MB) Additional Information: full citation, references, citings, index terms

49 The design of an interactive online help desk in the Alexandria Digital Library Robert Prince, Jianwen Su, Hong Tang, Yonggang Zhao

March 1999 ACM SIGSOFT Software Engineering Notes, Proceedings of the

international joint conference on Work activities coordination and collaboration, Volume 24 Issue 2

Full text available: pdf(1.53 MB)

Additional Information: full citation, abstract, references, citings, index <u>terms</u>

In large software systems such as digital libraries, electronic commerce applications, and customer support systems, the user interface and system are often complex and difficult to navigate. It is necessary to provide users with interactive online support to help users learn how to effectively use these applications. Such online help facilities can include providing tutorials and animated demonstrations, synchronized activities between users and system

supporting staff for real time instruction ...

Keywords: collaboration, digital library, online help desk, online support, user interface

50 Collaborative virtual workspace

Peter J. Spellman, Jane N. Mosier, Lucy M. Deus, Jay A. Carlson

November 1997 Proceedings of the international ACM SIGGROUP conference on Supporting group work: the integration challenge: the integration challenge

Full text available: pdf(1.03 MB)

Additional Information: full citation, references, citings, index terms

Keywords: MOO, MUD, collaboration framework, virtual environments

51 Disconnected operation in the Coda file system

James J. Kistler, M. Satyanarayanan

September 1991 ACM SIGOPS Operating Systems Review, Proceedings of the thirteenth ACM symposium on Operating systems principles, Volume 25 Issue 5

Full text available: (1.39 MB)

Additional Information: full citation, abstract, references, citings, index terms

Disconnected operation is a mode of operation that enables a client to continue accessing critical data during temporary failures of a shared data repository. An important, though not exclusive, application of disconnected operation is in supporting portable computers. In this paper, we show that disconnected operation is feasible, efficient and usable by describing its design and implementation in the Coda File System. The central idea behind our work is that caching of data, now ...

52 Extended ephemeral logging: log storage management for applications with long lived transactions



John S. Keen, William J. Dally

March 1997 ACM Transactions on Database Systems (TODS), Volume 22 Issue 1

Full text available: cdf(566.34 KB) Additional Information: full citation, references, index terms, review

Keywords: OLTP, disk management, logging, long transactions

53 Towards effective and efficient free space management

Mark L. McAuliffe, Michael J. Carey, Marvin H. Solomon

June 1996 ACM SIGMOD Record, Proceedings of the 1996 ACM SIGMOD international conference on Management of data, Volume 25 Issue 2

Full text available: pdf(1.34 MB)

Additional Information: full citation, abstract, references, citings, index

An important problem faced by many database management systems is the "online object placement problem"--the problem of choosing a disk page to hold a newly allocated object. In the absence of clustering criteria, the goal is to maximize storage utilization. For mainmemory based systems, simple heuristics exist that provide reasonable space utilization in the worst case and excellent utilization in typical cases. However, the storage management problem for databases includes significant additio ...

54 An analysis of database workload performance on simultaneous multithreaded processors



Jack L. Lo, Luiz André Barroso, Susan J. Eggers, Kourosh Gharachorloo, Henry M. Levy, Sujay S. Parekh

April 1998 ACM SIGARCH Computer Architecture News, Proceedings of the 25th annual international symposium on Computer architecture, Volume 26 Issue 3

Full text available: Additional Information: full citation, abstract, references, citings, index

Simultaneous multithreading (SMT) is an architectural technique in which the processor issues multiple instructions from multiple threads each cycle. While SMT has been shown to be effective on scientific workloads, its performance on database systems is still an open question. In particular, database systems have poor cache performance, and the addition of multithreading has the potential to exacerbate cache conflicts. This paper examines database performance on SMT processors using traces of th ...

55 Applications thrive on open systems standards

Eric J. Strand, Rajiv P. Mehta, Raju Jairam September 1994 StandardView, Volume 2 Issue 3

Full text available: pdf(890.71 KB) Additional Information: full citation, pitings, index terms

56 Designing database interfaces with DBface

Roger King, Michael Novak

April 1993 ACM Transactions on Information Systems (TOIS), Volume 11 Issue 2

Full text available: pdf(1.86 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

DBface is a toolkit for designing interfaces to object-oriented databases. It provides users with a set of tools for building custom interfaces with minimal programming. This is accomplished by combining techniques from User Interface Management Systems (UIMS) with a built-in knowledge about the specific kinds of techniques used by object-oriented databases. DBface allows users to create graphical constructs and interactive techniques by taking advantage of an object-oriented database envir ...

Keywords: graphical interfaces, object-oriented databases, user interface management systems

57 An object-oriented approach to data management: why design databases need it S. Heiler, U. Dayal, J. Orenstein, S. Radke-Sproull

October 1987 Proceedings of the 24th ACM/IEEE conference on Design automation

Additional Information: full citation, abstract, references, citings, index Full text available: xif(702.96 KB)

An object-oriented approach to management of engineering design data requires object persistence, object-specific rules for concurrency control and recovery, views, complex objects and derived data, and specialized treatment of operations, constraints, relationships and type descriptions. We discuss object-orientation as more than an implementation paradigm, and show how an object-oriented approach simplifies both use and implementation of engineering design systems.

58 Domain-independent natural language interfaces: Problems in natural-language interface to DBMS with examples from EUFID Marjorie Templeton, John Burger



February 1983 Proceedings of the first conference on Applied natural language processing

Full text available:

Publisher Site

Additional Information: full citation, abstract, references, citings

For five years the End-User Friendly Interface to Data management (EUFID) project team at System Development Corporation worked on the design and implementation of a Natural-Language Interface (NLI) system that was to be independent of both the application and the database management system. In this paper we describe application, natural-language and database management problems involved in NLI development, with specific reference to the EUFID system as an example.

59 Efficient data-parallel files via automatic mode detection

Jason A. Moore, Philip J. Hatcher, Michael J. Quinn

May 1996 Proceedings of the fourth workshop on I/O in parallel and distributed systems: part of the federated computing research conference

Full text available: Tool(1.34 MB)

Additional Information: full citation, references, citings, index terms

60 Development of an OO infrastructure for mainframe database applications

Darryl James Rothering

October 1994 ACM SIGPLAN Notices, Proceedings of the ninth annual conference on Object-oriented programming systems, language, and applications, Volume

Full text available: diff(839.77 KB) Additional Information: full citation, abstract, references, index terms

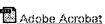
Large mainframe installations need and want to exploit the advantages of Object Technology (OT), but without totally abandoning their legacy environments. Implementing Object Orientation in such a COBOL/CICS/DB2 environment is a challenge: there is neither language support, nor development tools, nor execution infrastructure, nor testing utilities. Yet Object Orientation can be fully implemented, and a project can still meet rigorous performance requirements and tough delivery time scales. ...

Results 41 - 60 of 68

Result page: previous 1 2 3 4

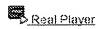
next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us











Search: The ACM Digital Library

The Guide

US Patent & Trademark Office

+metadata +location +lock +portion

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Published before March 1999 Terms used metadata location lock portion

Found 68 of 94,677

Sort results by

Display

results

relevance

expanded form

Save results to a Binder 2 Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

Results 61 - 68 of 68

Result page: previous 1 2 3 4

Relevance scale

61 Response to "Problems with DCE security services"

Walter Tuvell

April 1996 ACM SIGCOMM Computer Communication Review, Volume 26 Issue 2

Full text available: pdf(1.01 MB) Additional Information: full citation, index terms

window

62 An Extended Data-Flow Architecture for Data Analysis and Visualization

Greg Abram, Lloyd Treinish

October 1995 Proceedings of the 6th conference on Visualization '95

Publisher Site

Full text available: Additional Information: full citation, abstract, citings

Modular visualization environments utilizing a data-flow execution model have become quite popular in recent years, especially those that incorporate visual programming tools. However, simplistic implementations of such an execution model are quite limited when applied to problems of realistic complexity, which negate the intuitive advantage of dataflow systems. This situation can be resolved by extending the execution model to incorporate a more complete and efficient programming infrastructur ...

Keywords: visualization systems, data-flow, visual programming, graph analysis

63 APPL/A: a language for software process programming

Stanley M. Sutton, Dennis Heimbigner, Leon J. Osterweil

July 1995 ACM Transactions on Software Engineering and Methodology (TOSEM), Volume 4 Issue 3

Full text available: mindf(4 89 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Software process programming is the coding of software processes in executable programming languages. Process programming offers many potential benefits, but their realization has been hampered by a lack of experience in the design and use of process programming languages. APPL/A is a prototype software process programming language developed to help gain this experience. It is intended for the coding of programs to represent and support software processes including process, product, and p ...

Keywords: consistency management, multiparadigm programming languages, software process programming, transaction management

64 Structured document handling—a case for integrating databases and information retrieval



Klemens Böhm, Adrian Múller, Erich Neuhold

November 1994 Proceedings of the third international conference on Information and knowledge management

Full text available: Dxif(1.01 MB)

Additional Information: full cliation, abstract, references, index terms

In this paper we discuss the structured multimedia documents that will be, or already are, to some degree the communication backbone of the so-called superhighways. It will be shown that storage and retrieval of such documents will best be handled by an integration of database and information retrieval technologies. We assume documents to be structured with the help of standards like SGML/HyTime and represented by the multitude of formats currently used for multimedia data. Starti ...

65 UniSQL's next-generation object-relational database management system.



Albert D'Andrea, Phil Janus

September 1996 ACM SIGMOD Record, Volume 25 Issue 3

Object-Relational DBMSs have been receiving a great deal of attention from industry analysts and press as the next generation of database management systems. The motivation for a next generation DBMS is driven by the reality of shortened business cycles. This dynamic environment demands fast, cost-effective, time-to-market of new or modified business processes, services, and products. To support this important business need, the next generation DBMS must: 1. leverage the la ...

66 Semantic video indexing: approach and issues



March 1999 ACM SIGMOD Record, Volume 28 Issue 1

Full text available: 📆 pcif(634.24 KB) Additional Information: full citation, abstract, index terms

Providing concept level access to video data requires, video management systems tailored to the domain of the data. Effective indexing and retrieval for high-level access mandates the use of domain knowledge. This paper proposes an approach based on the use of knowledge models to building domain specific video information systems. The key issues in such systems are identified and discussed.

67 Query processing techniques in the summary-table-by-example database query language



Gultekin Özsoyoğlu, Victor Matos, Meral Özsoyoğlu

December 1989 ACM Transactions on Database Systems (TODS), Volume 14 Issue 4

Full text available: pdf(3.52 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Summary-Table-by-Example (STBE) is a graphical language suitable for statistical database applications. STBE gueries have a hierarchical subquery structure and manipulate summary tables and relations with set-valued attributes. The hierarchical arrangement of STBE queries naturally implies a tuple-by-tuple subquery evaluation strategy (similar to the nested loops join implementation technique) which may not be the best query processing strategy. In this paper we discuss the query ...

68 Contexts and metamessages in object-oriented database programming language

design

Michael Caruso, Edward Sciore



Full text available: pdf(1.02 MB)

Additional Information: full citation, abstract, references, citings, index iems

VISION is an object-oriented database system currently used commercially to develop investment analysis and other large statistical applications. Characteristic of these applications, beside the standard issues of structural and computational richness, is the need to handle time, versions, and concurrency control in a manner that does not produce combinatoric complexity in object protocol. This paper describes the approach taken by VISION in addressing these issues.

Results 61 - 68 of 68

Result page: previous 1 2 3 4

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2005 ACM, Inc. Terms of Usage Privacy Policy Code of Ethics Contact Us





